arXiv:1604.03458v1 [math.OC] 12 Apr 2016

Resource Allocation with Population Dynamics

Jonathan Epperlein and Jakub Mareček*

Abstract

Many analyses of resource-allocation problems employ simplistic models of the population. Using the example of a resource-allocation problem of Mareček et al. [Int. J. Control 88(10), 2015], we introduce rather a general behavioural model, where the evolution of a heterogeneous population of agents is governed by a Markov chain. Still, we are able to show that the distribution of agents across resources converges in distribution, for suitable means of information provision, under certain assumptions. The model and proof techniques may have wider applicability.

1 Introduction

There are resource-allocation problems encountered in almost every aspect of human lives: from utilities such as power systems and water systems, to transportation, and office space allocation. Many analyses of resource-allocation problems employ simplistic models of the population which ignore much of the complexity of human behaviour. Notice, for example, that the demand for a resource is often non-stationary, as exemplified by the work-day morning rush hour in transportation and the existence of predictable peaks in the demand in many other domains. Notice, further, that humans may have access to only very limited amount of information, but may still consider multiple criteria, and that their appreciation of the criteria may vary over time. As an example of a particular resource-allocation problem, we introduce a model of behaviour and the related demand process, which captures both the multi-criteria aspects of the decision making and non-stationarity of the demand process. Still, we show that the distribution of agents across resources converges in distribution, for suitable means of information provision and under certain assumptions.

As our running example, we consider the problems faced by transportation authorities in charge of a road network composed of a number of road segments. For each road segment, the travel time is, in principle, a time series with a data point per a vehicle passing across the road segment. Transportation authorities increasingly have information on traffic conditions throughout the road network

*J. Epperlein and J. Mareček are at IBM Research – Ireland, B3 IBM Campus Damastown, Dublin 15, Ireland.

at this level of detail, but want to broadcast much less information to the public. Mareček et al. [17, 18, 19] have recently demonstrated that such information provision should be seen as means of modulating the demand process and that different means of modulation may lead to very different outcomes. Whereas, for example, it is natural for each user to forecast the travel time as a scalar for each route, if the transportation authorities announce two distinct scalar values for two distinct alternative routes, the route with the lower scalar value announced may become congested, and if the authorities are truthful in reporting the congestion, the congestion may alternate between the two alternative routes in a sub-optimal limit-cycle behaviour, *ad infinitum*. In contrast, providing two intervals for two distinct alternative routes may benefit from the uncertainty in the response of the population. Mareček et al. used a rather simple behavioural model, where the response of each agent is determined by its level of risk aversion [19, 17] and possibly actuation delay [17], and the distribution of levels of risk aversion in the population is sampled in a memory-less fashion.

In Section 2, we introduce a considerably more realistic behavioural model, where the evolution of a heterogeneous population of agents is governed by a Markov chain. Section 3 gives a brief summary of relevant results in the theory of iterated function systems and their extension to recurrent iterated function systems; using these results, we show in Section 4 that the distribution of agents across resources converges in distribution, under certain assumptions. Section 5 illustrates the model and main result with a few simulations, and we explain the importance of the result in the context of related work in Section 6, before concluding with a variety of suggestions for future work in the field. Appendix A details the crucial step of generating a transition matrix for the Markov chain governing the population dynamics.

2 The Model

Let us consider a model consisting of three components [19, 17]: a set of resources, a central authority that provides information about the resources, and a population of agents, who make decisions as to which resource to use, based solely on the information made available by the central authority and their personal cost functions. Further, let us extend this model with population dynamics in Section 2.3.1. A summary of the model is presented in Figure 1. In the following, we describe each of the components in detail.

2.1 Resources

There is a finite set $\{1, 2, \ldots, M\}$ of resources, from which each agent chooses exactly one at every time step. For example in transportation, the resources would correspond to routes from origins to destinations provided by a road infrastructure. The cost of using such a resource is modelled by continuous

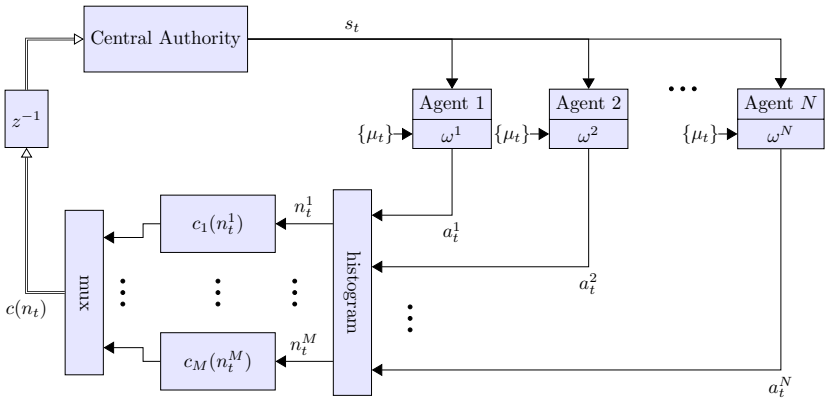

Figure 1: A signal flow of the model described in Section 2: The central authority broadcasts at time t the signal s_t based on information collected up until then (which does not include t itself). Every agent i makes its decision a_t^i which resource to use based on its policy ω^i. The decisions of all agents generate the new congestion profile n_t, which then leads to the current cost profile $c(n_t)$; z^{-1} denotes a unit delay.

cost functions ("link performance functions") of the form $c_m(n_t^m)$, where the index m denotes the resource, and n_t^m denotes the number of agents choosing resource m at time step t. Implicit in this form is the assumption of separability: the performance of resource m depends only on the utilization of resource m itself. In the transportation applications, the cost would correspond to the travel time of a single driver along the route. The vector $n_t \stackrel{\text{def}}{=} \begin{bmatrix} n_t^1 & \cdots & n_t^M \end{bmatrix}$ is the *congestion profile*.

2.2 Central Authority and Signalling Scheme

The central authority has access to the history of congestion profiles $\{n_\tau\}_{\tau=1}^{t-1}$ and provides information derived from this history to all agents, with the aim of lowering the social cost, $\sum_{m=1}^{M}(n_t^m/N)c_m(n_t^m)$, while being truthful, and hence trustworthy. Here, the same *two* pieces of information, u_t^m and v_t^m, are provided to all agents about *each* resource m. One can imagine u_t^m and v_t^m are, for example, the minimum and maximum costs within a time window, mean and variance over the same time window, or information about travel time and tolls at resource m; see Table 1 for details and references. In all cases listed, the signal $s_t = \begin{bmatrix} u_t^1 & v_t^1 & u_t^2 & \cdots & u_t^M & v_t^M \end{bmatrix}$ remains in $\mathbb{R}^{2M}$.

For this paper, we use exponential smoothing on the past costs of resource m

to obtain u_t^m, whereas v_t^m provides a measure of their volatility:

$$\begin{aligned} u_t^m &\stackrel{\text{def}}{=} (1-q_1)\Big(c_m(n_{t-1}^m) + q_1 c_m(n_{t-2}^m) + \cdots + q_1^{t-1} c_m(n_1^m)\Big) \\ v_t^m &\stackrel{\text{def}}{=} (1-q_2)\Big(\big|c_m(n_{t-1}^m) - u_{t-1}^m\big| + q_2\big|c_m(n_{t-2}^m) - u_{t-2}^m\big| + \cdots \\ &\qquad + q_2^{t-1}\big|c_m(n_1^m) - u_1^m\big|\Big), \end{aligned} \tag{1}$$

which can be implemented recursively as

$$\begin{aligned} u_t^m &= q_1 u_{t-1}^m + (1-q_1) c_m(n_{t-1}^m) \\ v_t^m &= q_2 v_{t-1}^m + (1-q_2)\,\big|c_m(n_{t-1}^m) - u_{t-1}^m\big|, \end{aligned} \tag{2}$$

thus requiring minimal storage at the central authority.

2.3 Agents and their Policies

Agents base their decisions at time t on the information provided by the central authority at time t by applying their policies. The heterogeneity of the agent population is taken into account by letting each agent have its own policy. Of course, a lot of policies are imaginable; we consider here a set of policies that is parametrized by $\omega \in [0,1]$, a weight in each agent's cost function: An agent using policy ω will choose the resource $m^\star$ that minimizes $\omega u_t^m + (1-\omega) v_t^m$. In other words, at time t agents with policy ω select resource

$$\pi^\omega(s_t) \stackrel{\text{def}}{=} \arg\min_{m=1,\ldots,M} \omega u_t^m + (1-\omega) v_t^m. \tag{3}$$

If for instance, u_t^m (resp. v_t^m) is the minimal (resp. maximal) cost of resource m observed in the last r time steps, then agents with $\omega = 1$ entirely disregard the worst case and could be said to be risk-seeking; similarly, $\omega = 0$ would be risk-averse and $\omega = 0.5$ risk-neutral. If instead a mean and toll are broadcast, ω parametrizes the trade-off between each agent's travel time and their wallet.

Of course, the population of agents in a road network changes over time. Mareček et al. [19] modelled this by assuming that 1) the number N of agents is fixed and 2) the distribution of policies ω among the agents is i.i.d., i.e. at each time step, the distribution is selected from a finite set of distributions $\{\eta_1, \ldots, \eta_K\}$ according to a distribution $d = (d_1, \ldots, d_K)$. A little more formally: at each t, choose μ_t from $\{\eta_1, \ldots, \eta_K\}$ according to $\text{prob}(\mu_t = \eta_k) = d_k$; the population of agents at time t is then μ_t.

2.3.1 Population Dynamics

Both assumptions – constant number of agents and i.i.d. replacement of the population at every t – are not very realistic. Here, we replace the assumption

of i.i.d. population renewal (according to the fixed distribution d) with a Markov chain that makes d dynamic; in this way the new population selected at each time step can depend on the previous population. For instance, it seems unlikely that a population of mainly risk-averse drivers at one time step will give rise to a population of mainly risk-seeking ones in the next instance, reflecting the fact that not all users leave the road network at once, or that there might be some correlation between, say, the time of day and what kind of drivers are on the road.

More precisely: As before, we assume that there is a family of possible distributions of the levels of risk aversion in the population, with a finite index set $\mathcal{K} = \{1, \dots, K\}$. Additionally, there is a Markov chain with K states and transition probability matrix $P \in [0,1]^{K \times K}$. The probability of appearance of population η_j, $j \in \mathcal{K}$, at iteration $t+1$ is now given by $\text{prob}(\mu_{t+1} = \eta_j \mid \mu_t = i) = p_{ij}$, i.e. the probability of a specific η_j depends on what the last observed population was. This setting is clearly more general: with $p_{ij} = 1/K \ \forall i, j \in \mathcal{K}$, we recover the original i.i.d. assumption.

The transition probabilities p_{ij} can be used to model different ideas about the mechanism behind the change in populations. Two straightforward ones are reviewed briefly here, more details are given in the appendix.

The Markov chain could be used to encode dependence of the population on the time of day. As a simple example, say that there is a morning, noon, evening and night population of drivers, perhaps many risk-seeking drivers at night, mostly risk-neutral commuters in the morning, etc. Say those populations for morning, noon, evening, and night are numbered as 1,2,3, and 4. The corresponding transition matrix and adjacency graph are:

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

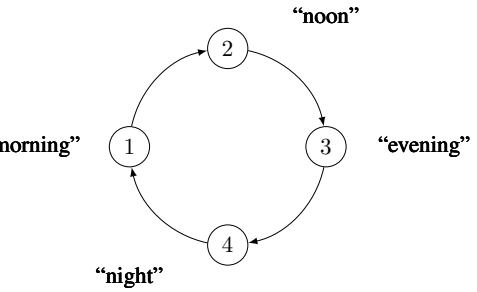

A simple generalization allowing for more than just one population per time of day can be achieved by replacing the "1"s by appropriate blocks, see Section A.1 for details.

If instead we are interested in a gradual change in population, e.g. due to individual drivers changing their behaviour or leaving and entering the road network, then the probability of the next population being η_j given that the current one is η_i should depend on "how different" they are from each other. In more

mathematical terms, let $\Delta(\cdot,\cdot)$ be a metric on the space of populations so that $\Delta(\eta_i, \eta_j)$ provides a notion of distance between the populations η_i and η_j. Then the probability p_{ij} should be a decreasing function of $\Delta(\eta_i, \eta_j)$: the farther apart η_i and η_j are, the less likely it is that one gives rise to the other in the next time step. To this end, we choose a parameter $\psi \in (0, 1)$ that reflects the probability of an agent changing its policy, and let p_{ij} proportional to $\psi^{\Delta(\eta_i,\eta_j)}$; see Section A.2.

In the case of drivers changing their policies, more drastic changes (e.g. going from risk-averse to risk-neutral) should contribute more to the distance than slight changes; this can be described by a discrete version of the *Wasserstein metric* (see Section A.2). If, however, we want to model drivers leaving the road network and being replaced, then the distance $\Delta(\eta_i, \eta_j)$ should depend only on how many drivers have to change policy to transform η_i into η_j (and not on the size of those policy changes); let us call this metric the *substitution metric*, since that is what it describes. Both of those notions are captured as special cases of the so-called Earth Mover's Distance (EMD, c.f. [27]). For details such as definitions and computation, see Section A.2.

3 Iterated Function Systems (IFS)

An *iterated function system* is a generalization of a Markov chain. It consists of a state space X with its metric d, a family W of Lipschitz functions[1] $W = \{w_j : X \to X \,|\, j \in \mathcal{K}\}$, where $\mathcal{K}$ is some index set (finite, countably infinite or worse, but let us assume here that it is countable) and a measure ν that makes $(\mathcal{K}, \cdot, \nu)$ a probability space.

At each iteration t of the IFS, j is selected from $\mathcal{K}$ according to ν and w_j is applied to the current state x_t to obtain x_{t+1}. Formally:

$$\text{prob}(X_{t+1} \in A | X_t = x_t) \stackrel{\text{def}}{=} \sum_{\mathcal{K}} \mathbb{1}_{\{i|w_i(x_t)\in A\}}(j)\nu(j),$$

i.e. the probability of x_{t+1} ending up in a set A is the probability of selecting an index j such that $w_j(x_t)$ is in A (the measure of the set of indices j for which $w_j(x_t)$ is in A). Here the Markov property is clear: the distribution of the next state X_{t+1} depends only on the current state x_t and not any "older" states x_{t-1} etc.

In this way, the IFS "jumps" around X. Unless we have a degenerate case such as all w_j having the same fixed point, we can not expect the sequence $\{x_t\}_{t=0}^{\infty}$ to converge in a classical sense; instead, we can have a weaker form, *convergence in distribution*: that there is a distribution Π on X such that as $n \to \infty$, the set $\{x_0, x_1, \ldots, x_n\}$ will be distributed according to Π.

[1] A function f on the metric space (X, d) is Lipschitz with constant s, or "s-Lipschitz," if for all $x, y \in X$, we have $d\big(f(x), f(y)\big) \leq s d(x, y)$.

Theorem 1 (E.g. [8, Thm. 1.1]). *Let L_j denote the Lipschitz constant of w_j and assume that the IFS is* contractive on average, *i.e.*

$$\sum_{\mathcal{K}} \nu(j) \log(L_j) < 0. \tag{4}$$

Then, there is a distribution Π on X such that $\{x_0, x_1, \ldots, x_n\}$ is distributed according to Π as $n \to \infty$.

If W is a family of contractions, i.e. if $L_j < 1$ for all j, then (4) is trivially satisfied.

IFS theory and congestion control The traffic model described above can be recast as an IFS, where the index set is $\mathcal{K} = \{1, \ldots, K\}$, and the randomly chosen function applied to the congestion profile n_t at each time step is parametrized by the population distribution η_j. Theorem 1 can then be applied to make the statement that there is a distribution Π such that, as $t \to \infty$, n_t follows this distribution. In other words, n_t will not be totally erratic, but it will be predictable in the sense that it will eventually behave like samples from a random variable with a fixed distribution.

3.1 Recurrent IFSs

The assumption in e.g. [19] of i.i.d. population renewal at each time step was necessary in order to apply Theorem 1, as the probability measure ν is not allowed to change over time or depend on the state x. However, extensions that relax this do exist, e.g. [20, 2]. Here, we consider the extension in [3], where a *recurrent iterated function system* (RIFS) is introduced as an IFS with an underlying Markov chain that modifies ν at each time step. More precisely:

We have an IFS as described in the last section with a finite index set $\mathcal{K}$, say $\mathcal{K} = \{1, \ldots, K\}$. Additionally, there is a Markov Chain with K states and transition probability matrix $P \in [0,1]^{K \times K}$. The probability of applying w_j at iteration $t+1$ is now given by $\text{prob}(i_{t+1} = j | i_t) = p_{i_t j}$, i.e. the probability of applying a specific w_j depends on what the last applied function w_{i_t} was! This is in contrast to the case of Section 3, where the probability to select a specific w_j was always the same and given by $\nu(j)$.

This setting is clearly more general: for a classical IFS, we can simply choose $p_{ij} = 1/K \ \forall i, j \in \mathcal{K}$. On the other hand, the way that X_t jumps around in X now is *not* a Markov process anymore — the distribution of X_t not only depends on X_{t-1}, but also on i_{t-1} — but the joint process of (X_t, i_t) jumping around in $X \times \mathcal{K}$ is.

Results analogous to Theorem 1 can be stated for this case, see e.g. [2, 3]. We state

Theorem 2 ([3]). *Assume we have an RIFS as described above, and let* $m : \{1, \dots, K\} \to [0, 1]$ *denote the stationary distribution of the underlying Markov chain (i.e.* m *corresponds to the normalized Perron eigenvector of* P^T*). Then, if*

$$\sum_{i=1}^{K} m(i) \log L_i = E_m\{\log L_i\} < 0, \tag{5}$$

there is a unique stationary distribution $\tilde{\nu}$ *of the Markov process* (X_t, i_t) *and* X_t *converges in distribution to* ν *with* $\nu(B) = \tilde{\nu}(B \times \mathcal{K})$*. Here,* L_i *again denotes the Lipschitz constant of* w_i*, and* E_m *denotes expected value with respect to* m*.*

Proof. This is just a corollary (much weaker, but sufficient for our purposes) to [3, Thm. 2.1 (ii)], which follows by taking $n = 1$ and removing the specifics of the stationary distributions. □ □

If (5) holds, we again say that the RIFS is *average contractive* or *contractive on average*.

That means that the model in Section 2 can be made more general while staying amenable to a similar stability analysis.

4 The Main Result

The traffic model described above can be recast as iterated applications of functions to the broadcast signals s_t, where the function to apply is chosen from a family whose index set is $\mathcal{K} = \{1, \dots, K\}$, with the next population of agents depending on the current population. At each iteration t, i_{t+1} is selected from $\mathcal{K}$ according to the i_t-th row of P, and the computation of the signals and the response of the population is captured by $w_{i_{t+1}}$ being applied to the current signal s_t to obtain s_{t+1}. That, and under which conditions, we can hope for convergence in distribution of the signalling process and the congestion profile is stated in the following theorem:

Theorem 3. *There exists a set of constants* κ'_m*,* $m = 1, \dots, M$ *such that if the costs* $c_m(\cdot)$ *are* $1/\kappa'_m$*-Lipschitz (with respect to the 1-norm), the signal* s_t *and the congestion profile* n_t *converge in distribution as* $t \to \infty$*.*

Proof.
Step 1: We begin by writing the signalling process $\{s_t\}$ generated by the model described in Section 2 and Figure 1 as an IFS on the state-space $X = \mathbb{R}_+^{2M}$ along similar lines as in [19]. Recall that

$$s_t = \begin{bmatrix} u_t^1 & v_t^1 & u_t^2 & \cdots & u_t^M & v_t^M \end{bmatrix}$$

and assume the signalling scheme described in (1) and (2); we repeat for convenience:

$$\begin{aligned} u_t^m &= q_1 u_{t-1}^m + (1-q_1) c_m(n_{t-1}^m) \\ v_t^m &= q_2 v_{t-1}^m + (1-q_2)\left|c_m(n_{t-1}^m) - u_{t-1}^m\right|. \end{aligned} \tag{6}$$

Now observe that

$$u_{t+1}^m = q_1 u_t^m + (1-q_1) c_m(n_t^m), \tag{7}$$

but

$$\begin{aligned} n_t^m &= \operatorname{card}\left(\left\{i \mid a_t^i = m\right\}\right) = \sum_{i=1}^N \delta_{a_t^i,m} = \sum_{i=1}^N \sum_{\omega\in\Omega} \delta_{\omega^i,\omega}\delta_{\pi^\omega(s_t),m} \\ &= \sum_{\omega\in\Omega} \underbrace{\left(\sum_{i=1}^N \delta_{\omega^i,\omega}\right)}_{\mu_t(\omega)} \mathbb{1}_{\{\sigma \mid \sigma u_t^m + (1-\sigma)v_t^m < \sigma u_t^s + (1-\sigma)v_t^s \ \forall s \neq m\}}(\omega) \\ &= \sum_{\omega\in\Omega} \mu_t(\omega) \quad \mathbb{1}_{\{\sigma \mid \sigma u_t^m + (1-\sigma)v_t^m < \sigma u_t^s + (1-\sigma)v_t^s \ \forall s \neq m\}}(\omega), \end{aligned} \tag{8}$$

where

$$\delta_{a_t^i,m} = \mathbb{1}_{\{j \mid a_t^j = m\}}(i) = \begin{cases} 1 & \text{if } a_t^i = m \\ 0 & \text{else.} \end{cases}$$

is the Kronecker symbol, so $\delta_{\omega^i,\omega}$ is 1 if agent i has policy ω, and $\delta_{\pi^\omega(s_t),m}$ is 1 if policy ω selects resource m. Now, (8) is unwieldy, but close inspection reveals that it shows that for fixed μ_t, n_t^m is just a function of s_t (through u_t^s and v_t^s, $s = 1, \ldots, N$), and plugging this into (6), we see that s_{t+1} can be written as a function of s_t, parametrized by the random variable μ_t:

$$s_{t+1} = w_j(s_t) \text{ if } \mu_t = \eta_j.$$

The size of the index set $\mathcal{K}$, i.e. the number K of possible distributions, is likely gigantic, namely

$$K = \binom{\operatorname{card}(\Omega) + N - 1}{N},$$

but nevertheless finite.

Step 2: Next, it will be established that if the c_m are Lipschitz, then so are the w_j of Step 1. For this, let $x, y \in \mathbb{R}_+^{2M}$, $x \neq y$ (they play the role of two possible s_{t-1}) and fix the population μ_t to some η_j. From (8) it then follows that the congestion profile at time t is just a function of the signal broadcast at s_{t-1}, i.e. x or y. Let us write $n_t^m(x)$ to reflect this. To investigate the Lipschitz constant of w_j, we write

$$\|w_j(x) - w_j(y)\|_1 = \sum_{m=1}^M R_1^m + R_2^m,$$

where

$$R_1^m \stackrel{\text{def}}{=} \left| q_1 x_{2m-1} + (1-q_1) c_m\big(n_t^m(x)\big) - \big[q_1 y_{2m-1} + (1-q_1) c_m\big(n_t^m(y)\big)\big] \right|$$

and

$$R_2^m \stackrel{\text{def}}{=} \Bigg| q_2 x_{2m} + (1-q_2)\left|c_m(n_{t-1}^m(x)) - x_{2m-1}\right| \\ - \Big[q_2 y_{2m} + (1-q_2)\left|c_m(n_t^m(y)) - y_{2m-1}\right|\Big] \Bigg|$$

where the first term in the sum corresponds to u_{t+1}^m and the second to v_{t+1}^m. We shall investigate those two terms separately, but first we repeat an important result from the proof of Theorem 1, specifically equation (15), in [19], namely that $n_t^m(\cdot)$ is Lipschitz in the 1-norm, so there is a constant κ_m such that

$$\left|n_t^m(x) - n_t^m(y)\right| \le \kappa_m \max_\omega \mu(\omega) \left\|x-y\right\|_1 \le \kappa_m N \left\|x-y\right\|_1 . \tag{9}$$

With (9), we then have

$$\left|c_m\big(n_t^m(x)\big) - c_m\big(n_t^m(y)\big)\right| \le L_m \left|n_t^m(x) - n_t^m(y)\right| \le L_m \kappa_m N \left\|x-y\right\|_1 ,$$

where L_m denotes the Lipschitz constant of c_m.

Now we can bound

$$R_1^m = \left|q_1(x_{2m-1} - y_{2m-1}) + (1-q_1) c_m\big(n_t^m(x)\big) - c_m\big(n_t^m(y)\big)\right| \le \\ q_1\left|x_{2m-1} - y_{2m-1}\right| + (1-q_1) L_m \kappa_m N \left\|x-y\right\|_1$$

and

$$R_2^m = \\ \Big|q_2(x_{2m} - y_{2m}) + (1-q_2)\big(\left|c_m(n_{t-1}^m(x)) - x_{2m-1}\right| - \left|c_m(n_{t-1}^m(y)) - y_{2m-1}\right|\big)\Big| \le \\ q_2\left|x_{2m} - y_{2m}\right| + (1-q_2)\big(\left|x_{2m-1} - y_{2m-1}\right| + \left|c_m(n_{t-1}^m(x)) - c_m(n_{t-1}^m(y))\right|\big) \le \\ q_2\left|x_{2m} - y_{2m}\right| + (1-q_2)\left|x_{2m-1} - y_{2m-1}\right| + (1-q_2) L_m \kappa_m N \left\|x-y\right\|_1 ,$$

where we used that $|a-b| - |a'-b'| \le |a-a'| + |b-b'|$, which can be derived using the reverse triangle inequality.[2] Reassembling and slight reordering of

[2] $|a-b| - |a'-b'| \le \Big||a-b| - |a'-b'|\Big| \le |a-b-(a'-b')| = |a-a'+b'-b| \le |a-a'| + |b-b'|$

terms yields

$$\|w_j(x) - w_j(y)\|_1 = \sum_{m=1}^{M} R_1^m + R_2^m \leq$$

$$\sum_{m=1}^{M} (1+q_1-q_2)\,|x_{2m-1} - y_{2m-1}| + q_2\,|x_{2m} - y_{2m}| + (2-q_1-q_2) L_m \kappa_m N\,\|x-y\|_1 \leq$$

$$\underbrace{\left(\max\{q_2, 1+q_1-q_2\} + (2-q_1-q_2) N \sum_{m=1}^{M} L_m \kappa_m\right)}_{\stackrel{\text{def}}{=} \bar{L}_\mu} \|x-y\|_1 \quad (10)$$

and so we have found an expression for (an upper bound of) the Lipschitz constant $\bar{L}_j$ of w_j.

Step 3: Now we are in a position to formulate sufficient conditions on L_m to allow application of Theorem 2. The condition

$$\sum_{j \in \mathcal{K}} m(j) \log \bar{L}_j < 0$$

is certainly satisfied independently of m (recall that m is the stationary distribution of the underlying Markov chain) if $\bar{L}_j < 1$ for all j. This is the case if

$$\sum_{m=1}^{M} L_m \kappa_m < \frac{1 - \max\{q_2, 1+q_1-q_2\}}{N\,(2-q_1-q_2)} \quad (11)$$

Hence the κ'_m postulated in the theorem statement need to satisfy

$$\sum_{m=1}^{M} \frac{\kappa_m}{\kappa'_m} < \frac{1 - \max\{q_2, 1+q_1-q_2\}}{N\,(2-q_1-q_2)}.$$

One choice is

$$\kappa'_m = \kappa_m M \frac{N\,(2-q_1-q_2)}{1 - \max\{q_2, 1+q_1-q_2\}}, \quad (12)$$

but there are obviously many more.

Hence, step 1 shows that the signalling process $\{s_t\}$ can be cast as an RIFS on the normed space $(\mathbb{R}^{2M}, \|\cdot\|_1)$ with a family of functions $\{w_j \mid j \in \mathcal{K}\}$, that by step 2 are Lipschitz. By step 3, under the right conditions, this RIFS is contractive on average and by Theorem 2, we then conclude that s_t converges in distribution. Since, by (8), n_t is a function of the two random variables s_t and μ_t, the congestion profile also converges in distribution. □ □

Remark 1. *Of course, Lipschitz constants have to be nonnegative, and thus* (11) *implies that in order for the proof to work, $q_2 > q_1$ has to hold.*

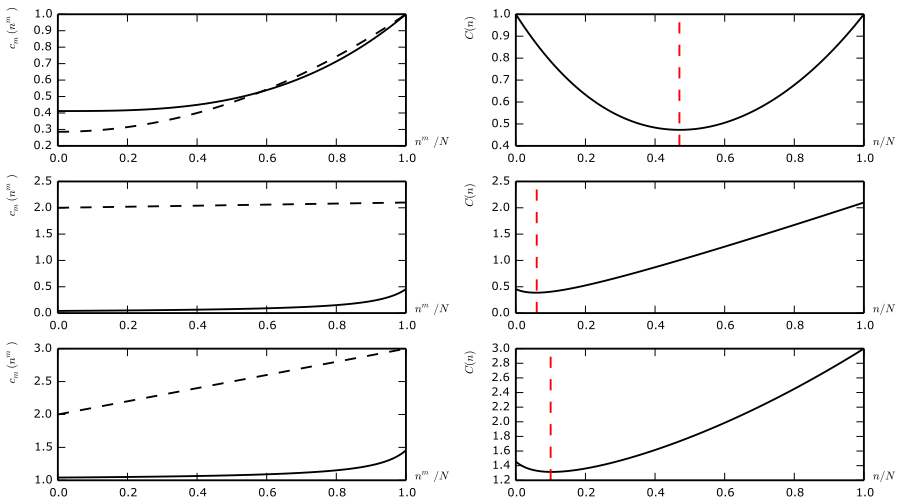

Figure 2: Three examples. Top left: Cost functions $x^2 + 0.4$ (dashed) and $(x^3 + 0.7)/1.7$ (solid line). Top right: The corresponding social cost with optimum of approximately 0.473 at 0.47 (dashed vertical line). Middle left: Cost functions $x/10+2$ (dashed) and $1+1/(1.1-x)/22$ (solid line). Middle right: The corresponding social cost with optimum of approximately 0.387 at 0.06 (dashed vertical line). Bottom left: Cost functions $x+2$ (dashed) and $1+1/(1.1-x)/22$ (solid line). Bottom right: The corresponding social cost with optimum of approximately 1.315 at 0.1 (dashed vertical line).

Remark 2. *While it is necessary for $c_m(\cdot)$ to be Lipschitz in order for the theory of RIFSs to apply, the bounds in* (12) *are far from tight, due to conservative approximations made in the proof. For instance in* (9)*, $\max_\omega \mu(\omega)$ is bounded by N. This bound is obviously only tight for populations where all agents have the same policy; for most populations, it is a gross overestimation. Additionally, numerical computations of m, the stationary distribution of the Markov chain, indicate that the stationary probability $m(i)$ of such populations, and hence their weight $m(i)$ in* (5)*, is typically very small, and this bound could most likely be improved upon.*

5 A Computational Illustration

Let us now illustrate the main result on a number of examples. Throughout, we use $M = 2$, three variants of $c_1(x)$ and $c_2(x)$ as suggested in Figure 2, and $q_1 = 0.45$, $q_2 = 0.5$. First, consider an example with $N = 2$ agents and $\Omega = \{0, 0.5, 1\}$, i.e., $\mathrm{card}(\Omega) = 3$ policies and $K = \binom{\mathrm{card}(\Omega)+N-1}{N} = 6$ possible distributions, assuming that every distribution possible in theory is also sensible in practice. One possible transition matrix for the case of the Wasserstein metric

is, up to rounding,

$$P_W = \begin{bmatrix} 0.35 & 0.18 & 0.12 & 0.12 & 0.07 & 0.06 \\ 0.22 & 0.28 & 0.18 & 0.18 & 0.11 & 0.09 \\ 0.14 & 0.18 & 0.29 & 0.12 & 0.18 & 0.14 \\ 0.14 & 0.18 & 0.12 & 0.29 & 0.18 & 0.14 \\ 0.09 & 0.11 & 0.18 & 0.18 & 0.28 & 0.22 \\ 0.06 & 0.07 & 0.12 & 0.12 & 0.18 & 0.35 \end{bmatrix},$$

whereas the corresponding transition matrix for the substitution metric is

$$P_S = \begin{bmatrix} 0.44 & 0.14 & 0.07 & 0.14 & 0.06 & 0.07 \\ 0.18 & 0.36 & 0.18 & 0.14 & 0.14 & 0.07 \\ 0.07 & 0.14 & 0.44 & 0.06 & 0.14 & 0.07 \\ 0.18 & 0.14 & 0.07 & 0.36 & 0.14 & 0.18 \\ 0.07 & 0.14 & 0.18 & 0.14 & 0.36 & 0.18 \\ 0.07 & 0.06 & 0.07 & 0.14 & 0.14 & 0.44 \end{bmatrix}.$$

For instance, this reflects that in the former case there is a 28-35% chance of the population remaining the same, whereas in the latter, this chance is 36-44%. For details on how these matrices are obtained, we again refer to Appendix A. The corresponding behaviour is captured in Figure 3, where each time-series has been obtained as the mean over 10,000 sample paths[3], with error bars at one standard deviation. Notice that in the case of $M = 2$, the social cost $\sum_{m=1}^{M}(n_t^m/N)c_m(n_t^m)$ is determined by n_t^1 alone, justifying the notation $C(n_t^1)$. Throughout, as proven, one observes the convergence in distribution.

Next, consider an example with $N = 20$ and $\Omega = \{0, 1/3, 2/3, 1\}$, i.e. 4 policies and 1771 distributions. The 1771×1771 transition matrices P_S and P_W are displayed in Figures 4 and 5. Note that generating each P_W requires solving $1770 \cdot 1771/2 = 1,567,335$ instances of the linear program (LP), as described in Appendix A. The corresponding behaviour is captured in Figure 6, again with each time-series obtained as the mean over 10,000 sample paths and error bars at one standard deviation. As in Figure 3, one observes the convergence in distribution throughout. Unlike in Figure 3, in the middle and bottom rows, one also observes convergence to a distribution concentrated around the optima of social cost of approximately 0.387 (middle) and 1.315 (bottom), respectively. Although the concentration of the limiting distribution around the true optimum cannot be proven generically, its appearance for larger N seems most intriguing.

6 Related Work

Our work builds upon a rich history of research in choice modelling [5, 26] and intelligent transportation systems [25]. There, any given user is associated

[3] Notice that although the horizon may seem short and the number of sample paths may seem low, this is justified by our main result, i.e., the convergence in distribution. Indeed, one may argue that convergence in distribution allows for a principled use of simulations.

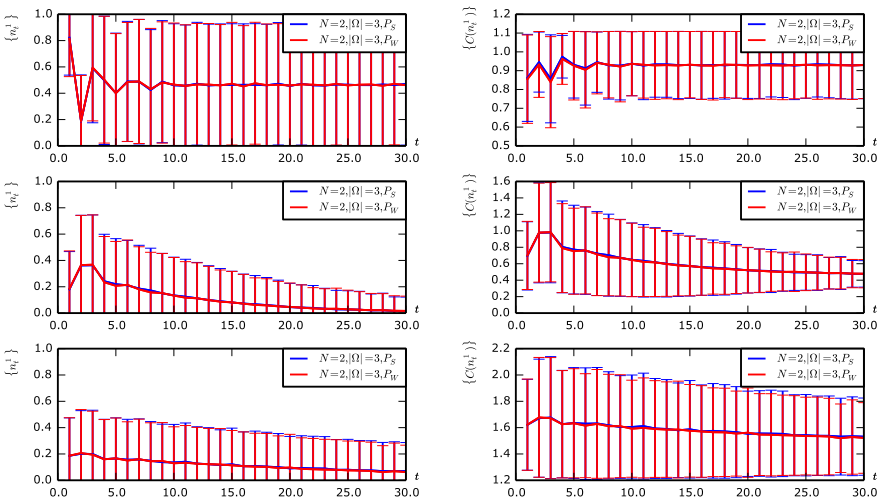

Figure 3: The behaviour of $N = 2$ agents on the three examples of Figure 2, considering both the Wasserstein metric (in red) and the discrete metric (in blue). Top left: Evolution of state over time for $x^2+0.4$ and $(x^3+0.7)/1.7$. Top right: The corresponding evolution of the social cost. Middle left: Evolution of state over time for $x/10+2$ and $1+1/(1.1-x)/22$. Middle right: The corresponding evolution of the social cost. Bottom left: Evolution of state over time for $x+2$ and $1+1/(1.1-x)/22$. Bottom right: The corresponding evolution of the social cost.

with an origin-destination pair, and hence with a set $\{1, 2, \ldots, M\}$ of routes connecting the two points. In the prevailing research direction, it is assumed that the perceived utility of a route m for user n is determined by a scalar $u_t^{n,m} \stackrel{\text{def}}{=} c_m(n_t^m) + \epsilon_t^{n,m}$, where $\epsilon_t^{n,m}$ is a random variable specific to the route, user, and time. The probability that user n picks $1 \leq m \leq M$ at time t is:

$$\begin{aligned} &\text{prob}_\theta(u_t^{n,m} < u_t^{n,m'} \ \forall\, m' \in \{1, 2, \ldots, M\} \setminus \{m\}) = \\ &\text{prob}_\theta(c_m(n_t^m) + \epsilon_t^{n,m} < c_{m'}(n_t^{m'}) + \epsilon_t^{n,m'} \forall\, m' \in \{1, 2, \ldots, M\} \setminus \{m\}), \end{aligned} \tag{13}$$

where θ is the probability distribution of the vector $\epsilon_t^n \stackrel{\text{def}}{=} (\epsilon_t^{n,1}, \epsilon_t^{n,2}, \ldots,)$. Clearly, the the next-step utilisation can then be expressed as an expectation of the perceived utility with respect to the distribution θ, or as by integrating over ϵ_t^n. There are a number of popular choices of θ, with varying complexity of the integration over ϵ_t^n.

The first such choice model [5] is the Multinomial Logit (MNL) of [9], which assumes that $\epsilon_t^{n,m}$ are independent and identically distributed Gumbel random variables. There, the equilibrium is computable as a convex program, under the assumption it exists and the cost functions are convex, but one may struggle to justify the fact that overlapping routes have uncorrelated $\epsilon_t^{n,1}$. In contrast, in

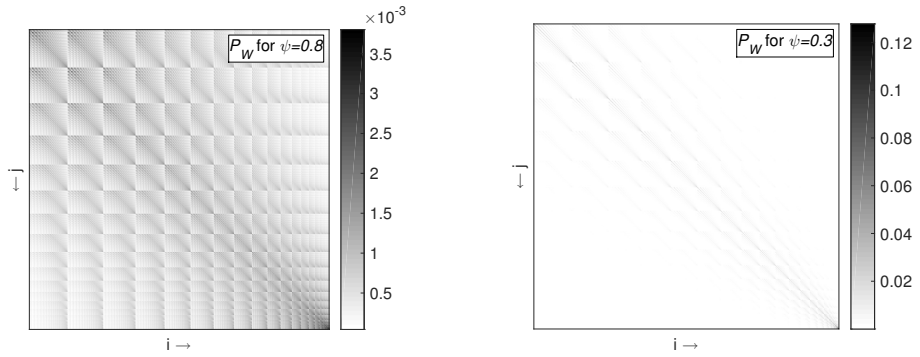

Figure 4: Transition matrices according to the Wasserstein metric with $N = 20$ agents and $\Omega = \{0, 1/3, 2/3, 1\}$ for two different values of ψ. The apparent fractal structure is an artefact of the order chosen for the populations η_i (Technically, this is referring to an order relation on the set $\{0, \dots, N\}^{\mathrm{card}(\Omega)}$, in which the possible populations make up the subset of tuples summing to N. The lexicographic order is chosen for the figures, however the effect of a different choice is merely cosmetic.)

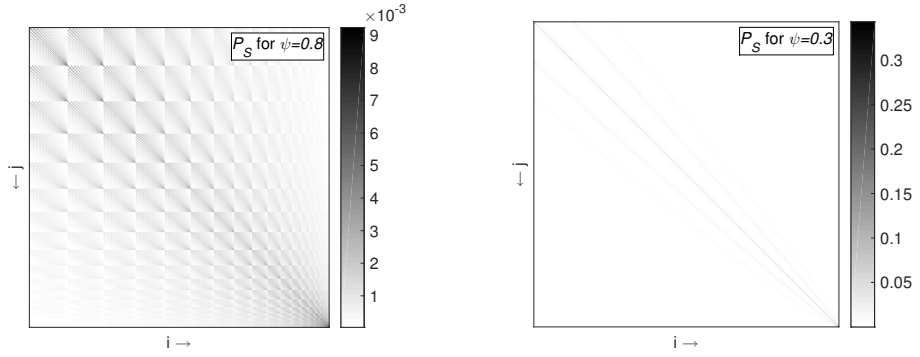

Figure 5: Transition matrices for the discrete metric with $N = 20$ agents and $\mathrm{card}(\Omega) = 4$ policies for two different values of ψ. The same comments as for Figure 4 apply.

the Multinomial Probit (MNP) model of [7], ϵ_t^n is assumed to be a multivariate normal random vector with mean 0 and a given covariance matrix. There, assuming the conditions for the existence of an equilibrium, one considers unconstrained minimisation of an integral in the computation of the equilibrium values. This has been extended in a number of ways, to obtain nested logit models [5], path-size logit [13], multinomial Weibit [6], c-logit [31], and yet more recent models [21, 1] considering distributionally robust optimisation over all distributions with mean 0 and the given covariance matrix. These models share a number of features, notably given travel costs $c_m(n_t^m)$ and stochastic $\epsilon_t^{n,m}$, and the existence of an equilibrium, under certain assumptions.

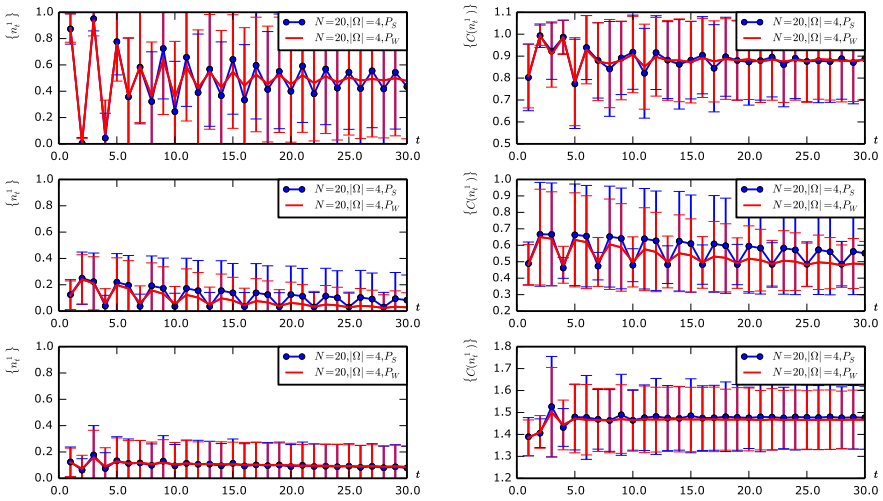

Figure 6: The behaviour of $N = 20$ agents on the three examples of Figure 2, considering both the Wasserstein metric (in red) and the discrete metric (in blue). Top left: Evolution of state over time for $x^2+0.4$ and $(x^3+0.7)/1.7$. Top right: The corresponding evolution of the social cost. Middle left: Evolution of state over time for $x/10 + 2$ and $1 + 1/(1.1 - x)/22$. Middle right: The corresponding evolution of the social cost. Bottom left: Evolution of state over time for $x+2$ and $1+1/(1.1-x)/22$. Bottom right: The corresponding evolution of the social cost. Notice the considerable difference in the evolutions of state and cost in the middle and bottom row, where $c_2(x)$ are the same and $c_1(x)$ are rather similar.

In an alternative line of research, one assumes that the travel costs $c_n(n_t^m)$ are a stationary random process, or a random variable with a distribution fixed over time. One can obtain a multi-parametric route choice model based on risk measures of the random variable and users' choice based on a level of risk-aversion weighing the risk measures, as in (3). Based on early simulations [15], a variety of such models have been analysed rigorously, recently, e.g., based on expectation and variance [4, 23] and expectation and value at risk [24, 22], which is sometimes known as the percentile equilibrium. Notice that in this line of research, the travel costs are exogenous random variable, which again makes it possible to guarantee the existence of an equilibrium, under certain assumptions.

Notice that much of the research on exogenous uncertainty in travel costs, as well as the research on choice models such as MNL and MNP, focuses on the conditions for the existence and computations of equilibria. In theory, the restriction to the existence of equilibria is rather unfortunate [11, 14], as it obscures the complexity of the possible behaviour of the non-linear dynamics. In practice, it may be very hard for the transportation authorities to steer the system to-

ward the equilibrium, while maintaining the trust of the users, which may be require the information provided to be truthful, in some sense. [17, 19] show that when one considers the truthful information provision, announcing a function of $c_m(n^m_{t-1})$ to all users at time t, the equilibria do not exists, in general. For example, in the simple case of truthful provision of costs $c_m(n^m_{t-1})$, with users' choice based on a scalar perceived utility, it is easy to observe a limit-cycle behaviour in a two-route example, where the traffic alternates between the routes, c.f., Appendix 1 of [19]. Even beyond this very simple control strategy, many standard controllers studied in control theory do fail [10] to stabilise the system. The study of the related dynamical systems without equilibria [30] is hence rather more recent, and often connected to the control of traffic lights (also known as signals), where it has been long recognised that the non-linear dynamics have limit cycles and yet more complex forms of behaviour.

Still, one may want to consider both the the costs $c_m(n^m_t)$ and their perceptions as non-stationary random processes, generated by the underlying non-linear dynamics. Mareček et al. [17, 19, 18] have studied measure-theoretic notions of convergence within such models, where $c_m(n^m_t)$ are non-stationary random processes, based on a closed-loop model with u^m_t and v^m_t being functions of the history of $n^m_{t-1}, n^m_{t-2}, \ldots$. Subsequently, the two-parametric choice model $\arg\min_{m=1,\ldots,M} \omega u^m_t + (1-\omega) v^m_t$ in (3), in the spirit of bounded rationality [28] is used, as in [4, 23, 24, 22]. An overview of the choices of parameters u^m_t and v^m_t is provided in Table 1. In (δ, γ)-signalling [17], there is one constant δ^m per resource m and one independent identically distributed (i.i.d.) uniform random variable ν^m_t with support $[-\delta^m/2, \delta^m/2]$ per resource m and time step t. Subsequently, one defines u^m_t and v^m_t as $c_m(n^m_{t-1}) + \nu^m_t - \delta^m/2$ and $c_m(n^m_{t-1}) + \nu^m_t + \delta^m/2$, respectively, resembling (13). In r-extreme signalling [19], (u^m_t, v^m_t) is the minimum and maximum over a time window of size r, i.e., within $n^m_{t-1}, n^m_{t-2}, \ldots, n^m_{t-r}$. Finally, in r-supported signalling [18], one minimises social cost C over the sub-intervals of minimum and maximum over a time window of size r, with certain restrictions on (u^m_t, v^m_t) captured by the inclusion in a set $P(S_t, \Omega)$, so as to minimise the next-step social cost. Many options remain to be explored.

7 Conclusions and Future Work

We have studied the stability of resource allocation, where a central authority broadcasts two values for each resource and agents take a convex combination of the two values, with the distribution of the convexifying coefficients across the population varying over time in a non-stationary fashion. Our contributions are as follows:

- A behavioural model, which does not simplify the nonlinear dynamics to a fixed point and considers evolution of the levels of risk-aversion in the population governed by a Markov chain.

- Novel means of information provision in resource-allocation problems, based on exponential smoothing of the past costs and past variance in the costs.
- A convergence result for the particular behavioural model and the particular means of information provision, whose proof technique may be more widely applicable.

This may spur much further research, along several directions.

One direction may focus on the wide variety of means of deriving the information to broadcast, some of which are suggested in Table 1. Outside of (A) further variants of the exponential smoothing, one could consider (B) the provision of the expected value and variance, c.f. [4, 23], (C) the expectation and value at risk for a given coefficient α and distribution function L with support $\{c_m(n_t^m)\}_t$, c.f. [24, 22], and (D) the expectation and conditional value at risk, also known as the expected shortfall, with the same distribution. For both the previously studied variants and those suggested in Table 1 for the first time, bounds on the concentration of the limiting distribution around its mean, and conditions for the mean being socially optimal, remain a major open problem. One may also bound the rate of convergence, c.f. [29], using the coupling argument of Hairer [12]. Although such arguments are somewhat technical, they allow for a principled use of simulations, with confidence bounds for a particular simulation set-up.

We also believe that the model of behaviour we have introduced is rich enough so as to allow the application in many other domains. For example in power systems, prices may be adjusted dynamically, but are currently not announced in real time, so as to prevent additional fluctuations in demand, leading to additional costs. A key open problem hence is how to provide the information on prices in a predictable and socially optimal fashion. One option may be to provide retail consumers with information on both the costs of energy and levels of pollution. Assuming each customer weighs these differently, one may apply a similar model for the evolution of the distribution of such a coefficient across the population. Similar reasoning may be applicable more broadly.

Acknowledgement Funding from the European Union Horizon 2020 Programme (Horizon2020/2014-2020) under grant agreement number 688380 is gratefully acknowledged.

Related work:

Ref.	Name	u_t^m	v_t^m
[17]	(δ,γ)-signalling	$c_m(n_{t-1}^m) + \nu_t^m - \delta^m/2$	$c_m(n_{t-1}^m) + \nu_t + \delta^m/2$
[19]	r-extreme signalling	$\arg\min_{j=t-r,\dots,t-1}\{c_m(n_j^m)\}$	$\arg\max_{j=t-r,\dots,t-1}\{c_m(n_j^m)\}$
[18]	r-supported signalling	$\text{proj}_{u_t^m} \arg\min_{(u_t^m,v_t^m)\in P(S_t,\Omega)} C(n_t)$	$\text{proj}_{v_t^m} \arg\min_{(u_t^m,v_t^m)\in P(S_t,\Omega)} C(n_t)$
Here	Exponential smoothing	$q_1 u_{t-1}^m + (1-q_1)c_m(n_{t-1}^m)$	$q_2 v_{t-1}^m + (1-q_2)\left\lvert c_m(n_{t-1}^m) - u_{t-1}^m\right\rvert$

Future work suggestions:

Ref.	Name	u_t^m	v_t^m
A	Variants of smoothing	$q_1 u_{t-1}^m + (1-q_1)c_m(n_{t-1}^m)$	$q_2 u_{t-1}^m + (1-q_2)c_m(n_{t-1}^m)$
B	Mean and variance	$\frac{1}{r}\sum_{j=t-r,\dots,t-1} c_m(n_j^m)$	$\frac{1}{r}\sum_{j=t-r,\dots,t-1}(c_m(n_j^m) - u_t^m)^2$
C	Mean and value at risk	$\frac{1}{r}\sum_{j=t-r,\dots,t-1} c_m(n_j^m)$	$\text{VaR}_\alpha \stackrel{\text{def}}{=} \inf\{l\in\mathbb{R} : \text{prob}(L>l)\le 1-\alpha\}$
D	Mean and conditional VaR	$\frac{1}{r}\sum_{j=t-r,\dots,t-1} c_m(n_j^m)$	$\text{CVaR}_\alpha \stackrel{\text{def}}{=} \frac{1}{\alpha}\int_0^\alpha \text{VaR}_\gamma d\gamma$

Table 1: A comprehensive overview of the related work and suggestions for future work within two-parameter information provision and two-parameter route choice formulations.

References

[1] Ahipasaoglu, S.D., Meskarian, R., Magnanti, T.L., Natarajan, K.: Beyond normality: A cross moment-stochastic user equilibrium model. Transport. Res. B: Meth. **81**(2), 333–354 (2015)

[2] Barnsley, M.F., Demko, S.G., Elton, J.H., Geronimo, J.S.: Invariant measures for Markov processes arising from iterated function systems with place-dependent probabilities. In: Annales de l'IHP Probabilités et statistiques, vol. 24, pp. 367–394 (1988)

[3] Barnsley, M.F., Elton, J.H., Hardin, D.P.: Recurrent iterated function systems. Constructive approximation **5**(1), 3–31 (1989)

[4] Bell, M.G., Cassir, C.: Risk-averse user equilibrium traffic assignment: an application of game theory. Transportation Research Part B: Methodological **36**(8), 671–681 (2002)

[5] Ben-Akiva, M.E., Lerman, S.R.: Discrete choice analysis: theory and application to travel demand, vol. 9. MIT press (1985)

[6] Castillo, E., Menéndez, J.M., Jiménez, P., Rivas, A.: Closed form expressions for choice probabilities in the Weibull case. Transportation Research Part B: Methodological **42**(4), 373–380 (2008)

[7] Daganzo, C.F., Sheffi, Y.: On stochastic models of traffic assignment. Transportat. Sci. **11**(3), 253–274 (1977)

[8] Diaconis, P., Freedman, D.: Iterated random functions. SIAM Review **41**(1), 45–76 (1999)

[9] Dial, R.: A Probabilistic Multipath Traffic Assignment Model which Obviates Path Enumeration. R & D report. Alan M. Voorhees and Associates for the United States Federal Highway Administration (1970)

[10] Fioravanti, A.R., Souza, M., Mareček, J., Shorten, R.N., Wirth, F.: On feedback design for smart cities. Working paper

[11] Haavelmo, T.: What can static equilibrium models tell us? Economic Inquiry **12**(1), 27–34 (1974)

[12] Hairer, M.: Coupling stochastic PDEs. In: XIVth International Congress on Mathematical Physics, pp. 281–289 (2005)

[13] Hoogendoorn-Lanser, S., van Nes, R., Bovy, P.: Path size modeling in multimodal route choice analysis. Transportation Research Record: Journal of the Transportation Research Board (1921), 27–34 (2005)

[14] Kerner, B.S.: Failure of classical traffic flow theories: Stochastic highway capacity and automatic driving. Physica A: Statistical Mechanics and its Applications (2016)

[15] Knoop, V.L., Bell, M.G., van Zuylen, H.J.: Traffic assignment based on individual risk-attitude. In: I.S. Society (ed.) Infrastructure Systems and Services: Building Networks for a Brighter Future (INFRA), 2008 First International Conference on, pp. 1–2. IEEE, Piscataway, NJ (2008)

[16] Luenberger, D.G., Ye, Y.: Linear and Nonlinear Programming. Springer, New York, NY (2008)

[17] Mareček, J., Shorten, R., Yu, J.Y.: Signaling and obfuscation for congestion control. Int. J. Control **88**(10), 2086–2096 (2015)

[18] Mareček, J., Shorten, R., Yu, J.Y.: Distributional robustness in congestion control. 11th European ITS Congress 2016 (2016)

[19] Mareček, J., Shorten, R., Yu, J.Y.: r-Extreme signalling for congestion control. Int. J. Control **to appear** (2016). DOI 10.1080/00207179.2016.1146968

[20] Mauldin, R.D., Szarek, T., Urbański, M.: Graph directed Markov systems on Hilbert spaces. In: Mathematical Proceedings of the Cambridge Philosophical Society, vol. 147, pp. 455–488. Cambridge Univ Press (2009)

[21] Mishra, V., Natarajan, K., Tao, H., Teo, C.P.: Choice prediction with semidefinite optimization when utilities are correlated. Automatic Control, IEEE Transactions on **57**(10), 2450–2463 (2012)

[22] Nie, Y.M.: Multi-class percentile user equilibrium with flow-dependent stochasticity. Transportation Research Part B: Methodological **45**(10), 1641–1659 (2011)

[23] Nikolova, E., Stier-Moses, N.E.: Stochastic selfish routing. In: Algorithmic Game Theory, pp. 314–325. Springer (2011)

[24] Ordóñez, F., Stier-Moses, N.E.: Wardrop equilibria with risk-averse users. Transportation Science **44**(1), 63–86 (2010)

[25] Papageorgiou, M., Ben-Akiva, M., Bottom, J., Bovy, P.H., Hoogendoorn, S., Hounsell, N.B., Kotsialos, A., McDonald, M.: ITS and traffic management. Handbooks in Operations Research and Management Science **14**, 715–774 (2007)

[26] Prato, C.G.: Route choice modeling: past, present and future research directions. Journal of Choice Modelling **2**(1), 65–100 (2009)

[27] Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. International Journal of Computer Vision **40**(2), 99–121 (2000)

[28] Simon, H.A.: Administrative behavior. Free Press, New York, NY (1957)

[29] Śleczka, M.: The rate of convergence for iterated function systems. Studia Math **205**(3), 201–214 (2011)

[30] Smith, M.: Dynamics of route choice and signal control in capacitated networks. Journal of Choice Modelling **4**(3), 30 – 51 (2011)

[31] Zhou, Z., Chen, A., Bekhor, S.: C-logit stochastic user equilibrium model: formulations and solution algorithm. Transportmetrica **8**(1), 17–41 (2012)

A The Transition Matrix P

In order to run simulations, and possibly also to obtain tighter theoretical results than Theorem 3, based on the model of population dynamics proposed in Section 2.3.1, it is necessary to compute a transition matrix P for the Markov chain generating the driver populations. We elaborate here on this procedure.

A.1 P to Encode Time Dependence

If P is used to model the dependence of the populations on an underlying temporal process, such as the time of day changing, the Markov chain would resemble a cycle: the state is forced to progress from morning to noon to afternoon and so on, until the day is over and everything repeats, there would be no transitions backwards. As explained in Section 2.3.1, the simplest encoding of that would yield a matrix P that is a unit shift, with ones on the first superdiagonal and a 1 one in the left bottom corner (if the different times of day are numbered appropriately). This would reflect a deterministic behaviour and thus again not be very realistic. Instead, there should be several possible populations per time of day. To illustrate the idea, consider the same example as in section 2.3.1, but assume there are now two different possible noon populations that both have the same probability. Then, P would be

$$P = \begin{bmatrix} 0 & .5 & .5 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix},$$

reflecting that if it was "morning" at time t, at time $t+1$ there is a 0.5 probability for each of the noon populations, but that in either case at $t+2$ we will have the afternoon population.

Formally: Assume there are T times of day, each with a set of N_ℓ different populations, all of which are equally probable. That would be reflected in a transition matrix

$$P = \left[\begin{array}{ccccc} 0_{N_1\times N_1} & \frac{1}{N_2}\mathbf{1}_{N_1\times N_2} & 0_{N_1\times N_3} & \cdots & 0_{N_1\times N_T} \\ \hdashline 0_{N_2\times N_1} & 0_{N_2\times N_2} & \frac{1}{N_3}\mathbf{1}_{N_2\times N_3} & \ddots & \vdots \\ \hdashline \vdots & & \ddots & \ddots & \\ \hdashline 0_{N_{T-1}\times N_1} & & & \ddots & \frac{1}{N_T}\mathbf{1}_{N_{T-1}\times N_T} \\ \hdashline \frac{1}{N_1}\mathbf{1}_{N_T\times N_1} & 0_{N_T\times N_2} & \cdots & & 0_{N_T\times N_T} \end{array}\right],$$

where $0_{M\times N}$ and $\mathbf{1}_{M\times N}$ denote $M \times N$ matrices of all zeros and ones, respectively.

A.2 Earth Mover's Distances (EMDs)

The Earth Mover's Distance (EMD) between histograms (which in the text have been also called populations and distributions) gets its name from interpreting a histogram as a pile of dirt. To transform one histogram into another, the dirt has to be rearranged; the *ground distance* between two to bins (or categories) is the work necessary to move one unit of dirt from one bin to the other, and the distance $\Delta(\eta, \gamma)$ between two histograms η and γ (with the same categories) is then defined as the minimal work necessary to transform one histogram into the other. For our purposes, the dirt corresponds to the agents, and the bins correspond to the policies Ω.

The choice of ground distance $h(\omega, \omega')$ between policies ω and ω' now depends on what brings about the change of individual policies. As already mentioned in Section 2.3.1, if an individual driver's policy does not change, but the driver instead leaves the road network and is replaced by a new driver, then $h(\omega, \omega')$ should only reflect that an agent has to be replaced, but be independent of ω and ω':

$$h(\omega, \omega') = 1 - \delta_{\omega,\omega'} = \begin{cases} 0 & \text{if } \omega = \omega' \\ 1 & \text{if } \omega \neq \omega'. \end{cases}$$

If instead the drift of an individual driver's policy is to be modelled, $h(\cdot, \cdot)$ should reflect the fact that drastic changes are less likely to occur than slight ones; in other words the ground distance $h(\omega, \omega')$ needs to be greater if the difference $|\omega - \omega'|$ is. One particularly simple choice is

$$h(\omega, \omega') = |\omega - \omega'| .$$

This special case of an EMD is better known – especially in the case of continuous distributions – as the Wasserstein distance, and we will call it by this name, too.

In either case, computation of $\Delta(\cdot, \cdot)$ requires solving an optimization problem, as there are many ways to "rearrange the dirt," but not all of them are optimal. For any choice of $h(\cdot, \cdot)$, the EMD can be computed as the solution of a linear program (LP), and if, as in our case, the items counted by the histograms cannot be divided into smaller units, an integer linear program (ILP): Let x_{ij} denote the amount of agents changing policy from ω_i to ω_j, and let the $\mathrm{card}(\Omega) \times \mathrm{card}(\Omega)$ matrix H have elements $H_{ij} = h(\omega_i, \omega_j)$. Then the minimal amount of work necessary to transform η into γ is the optimal cost of

$$\begin{aligned} &\text{minimize} && y = \mathrm{tr}(H^T X) = \sum_{i=1}^{M} \sum_{j=1}^{M} x_{ij} h(\omega_i, \omega_j) \\ &\text{subject to} && \begin{cases} x_{ij} \geq 0 \qquad x_{ij} \in \mathbb{Z} \\ \mathbf{1}^T X = \eta^T \\ X\mathbf{1} = \gamma. \end{cases} \end{aligned} \tag{LP}$$

Then,

$$\Delta(\eta, \gamma) \stackrel{\text{def}}{=} y^{\text{opt}}.$$

Remark 3. *For the substitution metric, we have*

$$H = \mathbf{1}_{\text{card}(\Omega)\times\text{card}(\Omega)} - I_{\text{card}(\Omega)}.$$

Remark 4. (LP) *has the form of a so-called transportation problem. Hence, if the problem data is integer, then there is always an integer optimum. In the present case, η and γ are integer by definition, so the only concern is H for the Wasserstein metric. However, if we choose all ω_i to be rational, then multiplying H by their least common denominator (or simply all denominators) renders H integer without changing the optimal solutions X^{opt}. This class of problems also allows more efficient solution than a general LP via special algorithms, see e.g. [16].*

The probability of transitioning from η to γ should of course be smaller if the distance $\Delta(\eta, \gamma)$ (i.e. the "work" required) is greater; hence, if now $\vec{\eta}_i$ and $\vec{\eta}_j$ denote two populations, then p_{ij} should be a bounded, nonnegative, decreasing function of $\Delta(\vec{\eta}_i, \vec{\eta}_j)$. One class of functions satisfying all requirements is $p_{ij} = \alpha\psi^{\Delta(\vec{\eta}_i,\vec{\eta}_j)}$, where $\psi \in (0, 1)$ and α is a normalization factor to make the resulting matrix P row-stochastic. Then, we can construct P from

$$\begin{aligned}\tilde{p}_{ij} &= \psi^{\Delta(\vec{\eta}_i,\vec{\eta}_j)} \\ p_{ij} &= \tilde{p}_{ij} \Big/ \left(\sum_k \tilde{p}_{ik}\right).\end{aligned}$$

The parameter ψ can be interpreted as a measure of the probability of one agent changing policy (for whatever reason). For $\psi = 1$, we get $p_{ij} = 1/K \ \forall i, j$, i.e. we recover the i.i.d. case of e.g. [19].